

# Parallel Matlab: RTEExpress on 64-bit SGI Altix with SCSL and MPT

*Cosmo Castellano*

Integrated Sensors, Inc.

Phone: 315-798-1377, x236

Email Address: [castellano@sensors.com](mailto:castellano@sensors.com)

## Abstract

By late 2004, RTEExpress™, a compiler and runtime environment that provides the capability for MATLAB® script files to be directly compiled and then executed on parallel high performance computers (HPC), will be released for the SGI platform, including the new Altix Itanium systems. [1] This new version of RTEExpress™ will take advantage of the SGI hardware and software package, specifically 64-bit operation, the SGI MPT (Message Passing Toolkit), the SGI SCSL (Scientific Computing Software Library), and is the first version of RTEExpress™ to utilize the advantages of a global shared-memory system. This paper presents the first test results using this new release. Improvement in corner-turn timing is anticipated due to the SGI NUMALink interconnect fabric, as compared to other interconnect technology common in Linux clusters, such as Ethernet. Up to an order of magnitude improvement in corner turn performance, and overall 2D FFT performance is expected.

## 1 Introduction

The RTEExpress™ environment is a software tool that assists a user in rapidly developing real-time embedded systems. RTEExpress™ is a compiler and runtime environment that provides the capability for MATLAB® script files to be directly compiled and then executed on parallel high performance computers (HPC). RTEExpress™ provides the capability to employ the power of an HPC on standard MATLAB® without having to recode the MATLAB® in the HPC target language. Its features include support for real-time data and machine performance visualization, multiple parallelization paradigms, multiple homogeneous parallel architectures, utilization of machine specific optimized vector libraries and native compilers, and the ability to change real-time algorithm parameters on-the-fly. [2]

The SGI Scientific Computing Software Library (SCSL) consists of several standard and proprietary scientific and math functions, optimized for use on the SGI platforms. Included in this package are BLAS (Basic Linear Algebra Subprograms) and LAPACK (Linear Algebra Package) libraries. The SCSL library

supports 64-bit integer arguments, single and double precision, and real and complex data types. [3] RTEExpress™ implementations have always taken advantage of vendor-supplied libraries, as possible, to fully exploit the target processing capabilities.

The SGI Message Passing Toolkit (MPT) combines the standard Message Passing Interface (MPI), which is utilized by RTEExpress™, with the SHMEM Library, which extends the interprocessor communication for shared memory systems. [4] The MPT facilities are a key element for taking full advantage of the SGI NUMALink Interconnect fabric. The Altix system combines the NUMAflex system architecture with the standard components, including the Intel Itanium 2 and the fully supported, 64-bit Linux operating system. [5]

The development of the RTEExpress™ environment was funded under DARPA/ITO BAA 95-19.

## 2 Performance Results

A MATLAB script performs the 2D complex FFT

```
init_matrix = ones(fftsize, fftsize) +
             j * ones(fftsize, fftsize)

loop
    store time t1
    a = fft(init_matrix)
    store time t2
    a = a'
    store time t3
    a = fft(a)
    store time t4
end loop
```

Elapsed times are computed, averaging time over several iterations.

The operation begins with initialization of a MATLAB matrix. Using RTEExpress, the script is mapped, using the graphic tool, "mapit", to several compute elements for a data-parallel operation. In this manner, the columns of the data matrix are distributed to the compute elements, thereby giving each compute element only a portion of the total number of columns to operate on. The resulting matrix is then transposed

<b>Report Documentation Page</b>			<i>Form Approved OMB No. 0704-0188</i>	
<p>Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p>				
1. REPORT DATE <b>01 FEB 2005</b>	2. REPORT TYPE <b>N/A</b>	3. DATES COVERED <b>-</b>		
4. TITLE AND SUBTITLE <b>Parallel Matlab: RTExpress on 64-bit SGI Altix with SCSL and MPT</b>			5a. CONTRACT NUMBER	
			5b. GRANT NUMBER	
			5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)			5d. PROJECT NUMBER	
			5e. TASK NUMBER	
			5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) <b>Integrated Sensors, Inc.</b>			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)	
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT <b>Approved for public release, distribution unlimited</b>				
13. SUPPLEMENTARY NOTES <b>See also ADM00001742, HPEC-7 Volume 1, Proceedings of the Eighth Annual High Performance Embedded Computing (HPEC) Workshops, 28-30 September 2004 Volume 1., The original document contains color images.</b>				
14. ABSTRACT				
15. SUBJECT TERMS				
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT <b>UU</b>	18. NUMBER OF PAGES <b>8</b>
a. REPORT <b>unclassified</b>	b. ABSTRACT <b>unclassified</b>	c. THIS PAGE <b>unclassified</b>		

and redistributed. This operation, the corner turn, is typically the limiting operation for 2D FFT performance. Lastly, the new columns are again

2D FFT timing. Systems utilizing Myrinet<sup>TM</sup> or DolphiNet<sup>TM</sup> interconnects show a notable improvement to standard 100base-T. The timings




Figure 1 – RTExpress 2D FFT Performance Timing

operated on by the compute elements, this time overwriting the input matrix. This in-place FFT should perform slightly faster since the step to copy the input data is not required.

The 2D FFT benchmark has been run on several platforms and interconnects, and Figure 1 shows some of the results taken for single-precision, complex 1k by 1k data set. Results have also been collected when running in double precision. The Corner turn performance (Fig 1b) provides an indication of the inter-processor communication capabilities of a particular system. To date, most systems have shown excellent scalable results for FFT performance (Fig 1a and Fig 1c), however, the total 2D FFT (Fig 1d) is limited by the performance of the corner turn. Most Linux clusters, utilizing standard 100base-T Ethernet have interconnect performance dominating the overall

from the SGI system will be compared to data displayed in Figure 1.

## References

- [1] M. Benincasa, R. Besler, D. Brassaw, R.L. Kohler, Jr. "Rapid Development of Real-Time Systems Using RTExpress", Proceedings of the 12<sup>th</sup> International Parallel Processing Symposium, pages 594-599, Mar 30 – Apr 3, 1998
- [2] SCSL (Scientific Computing Software Library) <http://www.sgi.com/software/scsl.html>
- [3] MPT (Message Passing Toolkit) <http://www.sgi.com/software/mpt>
- [4] SGI Altix 3000 <http://www.sgi.com/servers/altix>
- [5] M. Woodacre, D. Robb, D. Roe, K. Feind, "The SGI Altix 3000 Global Shared-Memory Architecture" SGI White Paper



*Parallel Matlab:  
RTExpress on 64-bit SGI Altix  
with SCSL and MPT*

**Cosmo Castellano**  
[castellano@sensors.com](mailto:castellano@sensors.com)

Integrated Sensors, Inc  
502 Court Street  
Suite 210  
Utica, NY 13502  
(315) 798-1377

[www.sensors.com](http://www.sensors.com)

# Matlab FFT Benchmark Test on Shared Memory SGI

---

- A Matlab script performs the 2D complex FFT

```
matrix = ones(fftsize, fftsize) + j * ones(fftsize, fftsize)
loop
```

```
    store time t1
```

```
    a = fft(init_matrix)
```

```
    store time t2
```

```
    a = a'
```

```
    store time t3
```

```
    a = fft(a)
```

```
    store time t4
```

```
end loop
```





fft1 elapsed time = t2 - t1

ctm elapsed time = t3 - t2





fft2 elapsed time = t4 - t3

- RTExpress is used to compile and run the MATLAB script on the 64-bit parallel computer on varying numbers of processors in data-parallel
- Elapsed times are computed, averaging time over several iterations
- First iteration is not counted

# Scaling: 64-Processor Altix



# SGI Shared Memory Improves Cornerturn and 2D FFT



Altix350 1.4 Ghz/3MB L3 cache

# Timing Information NOTES

---

- Please note that all timing information gathered is not intended to provide a recommendation for any particular hardware, but to illustrate parallel operation with various combinations of processors and interconnect systems
- Equipment used in the following tests may no longer be the hardware vendor's current offering
- “Improvement,” or speed-up, as compared to first-processor performance used to examine scaling rather than absolute timing
- Maximum RTExpress performance may be gained by fully using vector operations in MATLAB rather than using sequential loops

Parallel performance is extremely dependent on a specific application and implementation

# Parallel Matlab: RTExpress™ on 64-bit SGI Altix with SCSL and MPT

## What is RTExpress™

- Development and Runtime Environment allowing MATLAB scripts to be compiled and executed on real-time/parallel High Performance Computers (HPC)

– Provides a flexible means to harness the power of a HPC using MATLAB

– User does not require detailed knowledge of parallel programming

### – Supports:

- Embedded parallel architectures such as Mercury
- SUN Network of Workstations

• High performance Linux PC Servers

### – Support for FPGA functions

• Library of FPGA functions directly callable from MATLAB source

### • Now porting to SGI Altix systems

– [http://www.sgi.com/newsroom/press\\_releases/2004/june/altix\\_tcep.html](http://www.sgi.com/newsroom/press_releases/2004/june/altix_tcep.html)


– Intel / SGI Development Agreement

– Itanium 64-bit processing


– Shared Memory Architecture

– New RTExpress for SGI release expected by fall/winter of 2004

## Parallelization with RTExpress™ is Flexible and Efficient



## Developing Parallel Applications with RTExpress™



## Data Collection – 2D FFT Tests

SGI® Altix™ 3000



### • FFT Benchmark tests – Computation and Communication

#### • 1D FFT, transpose (cornerturn), 1D FFT in-place

#### • Testing on SGI Altix Linux servers with Shared Memory Interconnect

- Results of Shared Memory interconnect shows improved scaling on cornerturn

SGI® NUMAlink™ Interconnect Fabric



## 2D FFT Benchmark Test using RTExpress

### • A Matlab script performs the 2D complex FFT

```
matrix = ones(fftsize, fftsize) + j * ones(fftsize, fftsize)
loop
  store time t1
  a = f7init_matrix)
  store time t2
  a = a'
  store time t3
  a = fft(a)
  store time t4
end loop
```

• RTExpress is used to run the MATLAB script on varying numbers of processors in data-parallel

- Elapsed times are computed, averaging time over several iterations
- First iteration is not counted





• Please note that all timing information gathered is not intended to provide a recommendation for any particular hardware, but to illustrate parallel operation with various combinations of processors and interconnect systems

• Equipment used in the following tests may no longer be the hardware vendor's current offering

• Maximum RTExpress performance may be gained by fully using vector operations in MATLAB rather than using sequential loops

## 64-Processor System





Cornerturn shows excellent scaling up to Memory Bandwidth



64p 1.7GHz/9MB cache Altix

(note: production systems are 1.6GHz)

## Comparison to other Collected Data – SGI Shared Memory Improves Cornerturn



## Comparison to other Collected Data – SGI Shared Memory Improves Cornerturn

